

Amendments to the Specification:

Please replace the paragraph on page 22, lines 10-13 with the following paragraph that is between the two sets of dash symbols "--":

--Information about the BuDDy package can be obtained from Jorn Lind-Nielsen, Computer Systems Section at the Department of Information Technology, Technical University of Denmark, or from <http://www.itu.12/research/buddy/>, or by sending an e-mail to buddy@itu.12 who may be contacted by email or through the University website.--

Please replace the paragraph on page 27, lines 9-18 with the following paragraph that is between the two sets of dash symbols "--":

--In this pseudo code, T_0, i_0, i_max, alpha, beta are parameters for the simulated annealing algorithm. These parameters might be specified by a user or set by the optimizer. In some embodiments, alpha equals 0.98, and beta equals 1.1. The selection of the annealing parameters is well studied. One scheme for specifying these parameters is disclosed in "A Comparison of Annealing Techniques for Academic Course Scheduling," by M.A. Saleh Elmohamed, et al., published at 2nd international conference, PATAT97. *See also, e.g.,* "Simulated Annealing and Combinational Optimization," by Surendra Nahar, et al., University of Minnesota, 23 Design Automation Conference, pp. 293-299. Also, several software packages are available for determining the parameters for simulated annealing. One such package is ASA, written by Lester Ingber, www.ingber.com who may be contacted by email.--

Please replace the paragraph beginning at page 31, line 21 and ending on page 32, line 2, which starts with "After 610" with the following paragraph that is between the two sets of dash symbols "--":

--After 610, the query manager determines (at 615) whether the table manager returned any replacement sub-network for the set of indices it received at 610. If not, the process transitions to ~~660~~ 655, which will be explained below. Otherwise, the query manager selects (at 620) one replacement sub-network retrieved by the table manager at 610.--

Please replace the paragraph beginning at page 65, line 19 and ending on page 66, line 5, with the following paragraph that is between the two sets of dash symbols "--":

--Numerous known techniques can be used to construct all graphs for a given number of nodes and edges. Some embodiments construct initially all undirected graphs for a given number of nodes and edges. There are software packages available for constructing all undirected graphs. One such package is the "geng" program package by Brendan D. Mckay (bdm@cs.anu.edu.au) (who may be contacted by email), Computer Science Department, Australian National University. This package can be downloaded through the internet at:

<http://cs.anu.edu.au/~bdm/nauty/gtools10beta11.tar.gz>

After generating all undirected graphs, these embodiments generate all directed graphs by trying all possible assignments for directions on all edges of each graph. After constructing all directed graphs for each combination of nodes and edges, the process discards all cyclic graphs for the combination, and then stores each remaining graph in the graph table so long as the graph is not isomorphic to a previously stored graph.--

Please replace the paragraph on page 66, lines 6-13, with the following paragraph that is between the two sets of dash symbols "--":

--Checking for cycles and identifying isomorphic graphs is commonly known in the art. For instance, Cormen, Leiserson, Rivest and Stein, Introduction to Algorithms, Second Edition,

Chapter 22 (Elementary Graph Algorithms), MIT Press 2001 discloses one manner of checking a graph for cycles by traversing the graph. In addition, there are software packages available for identifying isomorphic graphs. One such package is the "nauty" package by Brendan D. Mckay (bdm@cs.anu.edu.au) (who may be contacted by email), Computer Science Department, Australian National University. This package can be downloaded through the internet, at:

<http://cs.anu.edu.au/people/bdm/nauty...>

Please replace the paragraph beginning at page 72, line 13 and ending on page 73, line 3, which starts with "After 2240" with the following paragraph that is between the two sets of dash symbols "--":

--After 2240, the process determines (at 2242) whether it has examined all the graphs generated at 2242-2202. If it has not examined all the graphs, it transitions to 2228 to select an unexamined graph. Otherwise, it discards (at 2244) multiple definitions of the same network or nearly the same networks. This is done by deleting all but one out of each group of generated sub-networks that have (1) the same graph table index, and (2) the same sets of function indices (defined at 2240). Such duplicate networks may appear for example because of symmetries of the graph structure. Based on the list of network table indices and function indices, the process then completes (at 2246) the database tables 525. Specifically, the process first creates the network table 1510, then the secondary index table 1505, and then the pivot-index table 1500. As mentioned above, the network table is sorted in an order specified by its stored network indices, the secondary table is sorted in an order specified by its stored primary and secondary indices, and the primary table 1500 is sorted in an order specified by its stored primary indices. After 2244-2246, the process ends.—